



Making Music

Computers help us create all kinds of music, but what of those that compose their own scores and interact with musicians? New breakthroughs in musical composition signal a turning point in the partnership of man and machine. **Jon Thompson** considers the impact of true computer creativity

Nearly every piece of recorded music passes through a number of computers before it reaches us. Computers are also used during composition, whether it's for dance music production or orchestral film scores. But how about computers that compose music automatically? Are we prepared for where a steady stream of breakthroughs might lead us?

The process of musical composition is largely one of following rules. Pop songs tend to have familiar layouts and, in classical music, forms such as the sonata have unmistakable structures. In fact, any genre or style of music can be described in terms of the rules it follows. One thing computers are very good at is following rules. By coding rules for making certain types of music, composers and scientists have been able to write programs that generate scores automatically and even mimic great maestros. This process of making so-called generative music has produced surprising results.

AUTOMATIC FOR THE PEOPLE

Mozart was one of the first people to think seriously about generating music automatically. In 1787, he published the *Musikalisches Würfelspiel*,

the musical dice game. The idea was deceptively simple. The game consisted of a book of short musical phrases and instructions on how to use a pair of dice to create a never-ending stream of unique minuets that sounded as if Mozart himself had composed them.

Mozart's game couldn't really produce unique music, though we now know his system had 1,328 possible combinations. This is still a high number and to non-expert ears, it really does seem like a constant stream of musical creativity. There's an online demonstration of the game at <http://sunsite.univie.ac.at/Mozart/dice>, where you can create, play and save your own minuet as a MIDI file.

As impressive as this musical dice game was, composing music remained a manual process until the middle of the 20th century and the invention of the digital computer.

MUSIC FROM CHAOS

One of the first to realise the potential computers held for writing

music was the influential Romanian-born Greek architect, mathematician and composer Iannis Xenakis. In the 1950s, Xenakis became interested in using randomness in guiding his compositions. Rather than leave everything to chance, however, Xenakis carefully formulated complex statistical rules covering how everything should fit together.

Because of the complexity of his way of working, the number of calculations needed to create his abstract soundscapes by hand became overwhelming. In the early 1960s, he turned to computers. In his 1962 book *Formalised Music* he discussed the process of using a PC to generate music by having, as he described it, "A list of note densities and probabilistic weights supplied by the programmer, leaving specific decisions to a random number generator".

That same year, he obtained access to a 7090 mainframe at IBM France and thus had the number-crunching power needed to try out his ideas fully. Writing in Fortran, he



An idea Mozart had for generating minuets automatically still influences modern research into computer-generated music



Xenakis was one of the first to realise that computers could aid him in writing complex, abstract music

began programming his rules for musical creativity and used the resultant software to produce the scores for what became known as his ST works. His programs helped him produce some of the most important works of modern classical music. In 1991, aged 79, he started on the culmination of his life's work, the program Gendyn. He wanted this to be a sort of black box capable of writing complete Xenakis-style works in its own right.

You can hear some examples of Xenakis' work at <http://tinyurl.com/7rdep>. Of course, musical appreciation is in the ear of the beholder; much depends on whether you think the work of modern composers sounds like heaven or hell. Still, there's no denying Xenakis's standing as one of the 20th century's leading composers.

NEW MUSIC, MAESTRO

While Xenakis used randomness, other composers have taken existing works as their software's starting point. One such composer is Professor David Cope of the University of California at Santa Cruz.

In 1981, Cope decided to write software that would help him overcome his writer's block by offering suggestions as to the next note or bar of music he should use when composing his works. The result was a program called Experiments in Musical Intelligence (EMI). To create this software, he had to code everything he knew about his own musical style so EMI could be called upon to generate meaningful suggestions. But uncovering and coding these subtle rules proved daunting, so he decided to code what he knew about the composition styles of famous composers such as Bach and Mozart to prove the concept worked.

After some trial-and-error programming of what Cope and other musicologists believed to be the mental processes composers used in creating their music, EMI began producing output. Though this sounded like variations of the real thing, he said, "The music seemed, at least to the educated ear, lifeless and without much musical energy." He needed EMI to learn from the maestros directly.

To do this, Cope revised EMI to examine databases containing the musical scores of target composers. He reasoned that existing works contain all the instructions needed to create new but related music based on how individual composers created the originals. He programmed EMI to analyse these originals and uncover the underlying instructions. To do this, he used a process called recombination, whose underlying concept is like that of Mozart's dice game.

Describing recombination Cope said, "All the great books in the English language are constructed from recombinations of the 26 letters of the alphabet. Similarly, most of the great works of

HOW TO... Make generative music with Koan Pro

Generative music is created by repeatedly processing a set of rules about what notes to play next. One software package designed to generate music this way is Koan Pro from UK software house SSEYO. Rather than creating random sequences of sounds, Koan Pro claims to be capable of producing everything from ambient sound landscapes to "beat-driven, floor-shaking dance".

Koan Pro is slightly non-deterministic. If you start with exactly the same parameters twice, the resultant music is somewhat different each time. As Koan Pro decides which notes to play next according to a set of user-definable rules, these provide it with options rather than a fixed sequence.

In 1996, pioneering producer and composer Brian Eno used an earlier version of the software to create the album *Generative Music 1*. You can hear examples of his work with Koan Pro at <http://tinyurl.com/cfp6s>.

The list of parameters you can use to control the direction Koan Pro takes is extensive. There's also an onboard synthesiser for those who want to generate their own sounds rather than rely on Windows' built-in software synthesiser. For the beginner, the range of options can seem intimidating, so we've put together a quick-start guide.





1 Download a demonstration version of the software from www.sseyo.com/products/koanpro. Click the Download link, save the file to your hard disk and run it to install the package. Run Koan Pro and create a new document. This brings up a small window allowing you to decide on a basic template. Select Backbeat and press OK. This creates a grid consisting of several instruments and various

parameters. Press Play and Koan Pro will begin a gradually changing rhythm.

2 Now add a funky bass part. We'll crossbreed the data from two existing instruments to do this. Stop the music, then right-click the Patch column entry for instrument 1 and select Mutate... from the resulting pop-up menu. Select Add Child from the submenu. This brings up a dialog box asking which of the existing instruments you'd like to use as the other parent. It also asks how much mutation it should introduce. A good level here is 25%. Select an instrument as the other parent then click OK to create your new instrument.

3 We need to change the sound of the new instrument to that of a bass guitar. Click the numbered, left-hand column entry for the new instrument to bring up the voice parameter dialog. In this window, first click on the Value column entry for the Patch parameter. Scroll up the resultant list of sounds until you reach 035-Electric Bass (pick). Change the MIDI channel to 0 by double-clicking on the relevant value and entering a new one. Press Play to hear your new bass part playing along to the rhythm.

4 You can also change the existing rhythm sounds to add elements such as pianos. Don't forget to change the MIDI channel to 0, though. Given the sheer number of parameters and rules you can manipulate and define, you should be able to produce some interesting, ever-changing music with Koan Pro in a range of styles. You can hear examples of what others have produced at www.sseyo.com/showcase.





Western art music exist as recombinations of the 12 pitches of the equal-tempered scale and their octave equivalents. The secret lies not in the invention of new letters or notes but in the subtlety and elegance of their recombination."

Cope reprogrammed EMI to analyse pieces of music to uncover the rules about which sequences of notes or phrases the original composer tended to use, and used these to create new works in a more intelligent form of Mozart's musical dice game. See our 'Chain reaction' box, right, for more details.

The music EMI produced became more pleasing, even to the trained ear. It managed to write a 'new' piece of Mozart (symphony No 42), and even fooled audiences into thinking that a human had created pieces that it had composed. To hear what EMI is capable of, visit <http://tinyurl.com/8gu9n>.


CHILDREN OF THE EVOLUTION

Xenakis and Cope made important breakthroughs, but true machine creativity has always been a Holy Grail. More recently, several promising approaches have taken their cue from the most creative process of all: nature itself.

At Sony's Computer Science Laboratory in Paris in 2002, for instance, Professor Eduardo Reck Miranda used genetic algorithms in a technique called evolutionary musicology to evolve a stream of new music. Genetic algorithms take a population of slightly varying datasets and crossbreed them to create a new generation. From these, they take those that best fit some predetermined criteria and breed again from them. By repeating this process, the datasets – which could be anything from the solution to an equation, to the design of a new turbine blade – gradually become better, as only the fittest parents get to breed in each generation.

In the case of an equation, the fittest formulae come up with answers nearest to the desired solution. As they evolve, the fittest of each generation of turbine blade design might be the most fuel-efficient. Once set in motion, this process is entirely blind, just as it is in nature, even though it creates seemingly designed results. Genetic algorithms, sometimes evolving solutions at top speed on supercomputers, are now used to generate and test successively better designs in industries such as car manufacture and chemical engineering.

Miranda's evolutionary musicology system consisted of a grid of programs, each running a genetic algorithm and communicating with their nearest neighbours. Starting randomly, each program swapped and crossbred new generations of musical ideas with those held by its nearest



Composer and producer Brian Eno used the generative music program Koan Pro to compose a complete album

CHAIN REACTION

How existing data can be used to create new output

One of the key ideas used in many forms of generative music, especially when emulating other composers, is the Markov chain. Named after Russian mathematician Andrey Markov, this is the result of analysing a stream of data and noting which values tend to follow which. The process of analysing the data counts the number of times that values A, B and C follow value X in the original stream.

The data can be anything, including words and musical notes. Software can construct new streams of data based on its analysis of what tends to follow what. To select successive values, rather than selecting purely at random, it uses the analysis of the number of times each A, B or C follows each X as an indication of how often it should choose each to follow every successive X.

The results of using Markov chains to generate text can be bizarre. Jo Wood of City University in London has written an online

Markov chain text generator (<http://tinyurl.com/9wffjb>). Entering a reasonably long piece of text (the longer the better) then pressing the Build Chain button followed by the Generate Text button creates new output. Entering a long speech, movie script or book you know well generates some seemingly intelligent and thoughtful output, as well as comical, nonsensical word associations. For some reason, the script to cult film favourite *Withnail & I* is particularly funny when processed in this way.

Creating Markov chains from text gives abstract results, as it doesn't follow the rules of grammar. But music lends itself well to the creation of Markov chains that tend to sound like variations on a theme. It makes you wonder how long it will be before someone decides to analyse the biggest hit singles ever to produce never-ending variations guaranteed to empty young pockets everywhere.

neighbours, which most closely resembled its own. Overall, the individual ideas generated by the group evolved as they ebbed and flowed across the grid. This process mimics the way in which individuals in a band take and modify ideas developed by other members to create something pleasing to the whole group – and, hopefully, the audience too.

Miranda fed the resultant ideas into a second program called ChaoSynth, which translated the data produced by the grid and played the resulting sounds through a speaker. To do this, ChaoSynth used a technique known as granular synthesis to create a smooth stream of sound corresponding to the program grid's blocky output states.

"This synthesis technique can be metaphorically compared with the functioning of a motion picture in which an impression of continuous movement is produced by displaying a sequence of slightly different images at a rate above the scanning capability of the eye," said Miranda. Because the grid's output evolved slowly, so did the output of the ChaoSynth, creating shifting, ambient soundscapes that had no human composer.

ALL THAT JAZZ

As impressive and creative as Miranda's work was, it was not toe-tapping stuff. For that, we must turn our attentions to the Rochester Institute of Technology and the Al Biles Virtual Quintet: a working jazz band with just one human member.

John 'Al' Biles is a jazz trumpeter and professor of information technology who regularly plays improvised concerts with a band consisting of himself and an Apple Powerbook. On the Powerbook is a program called Band in a Box, which produces an accompaniment for Biles and the other creative member of the ensemble. Called GenJam (short for genetic jammer), this software uses a microphone and sound-to-MIDI converter to translate Biles' trumpet playing into a stream of data to which it can react.

GenJam, written using Carnegie Mellon University's shareware MIDI Toolkit, has a pool of favourite musical ideas forming its basic style.

It codes these into what Biles calls chromosomes, which are pointers to complete musical phrases. Unlike those found in Mozart's dice game, however, GenJam breeds these phrases to produce new ones according to general rules about pitch and harmony, but also according to the ideas the system detects in Biles' playing. But does playing with GenJam feel like jamming with another human?

"[GenJam's] spontaneity and responsiveness is actually better than most of the humans I've jammed with," said Biles. "And my conversations with GenJam are often more musical than those with humans." Perhaps uniquely for software, the system can also react to audience feedback.

Before each performance, each member of the audience receives a paddle, like a table tennis bat, painted red on one side and green on the other. During a concert, they use these to indicate whether they like or dislike GenJam's output. An operator uses this visual feedback to influence GenJam to develop or discard ideas from its current pool as the concert progresses. This has had unexpected consequences for Biles, however. "The audience invariably greet the human's first solo with a blizzard of red and green," Biles explained. "Technology can be fun!"

Examples of GenJam and details of the album Biles produced using it can be found at <http://tinyurl.com/cb6px>.

NEVER MIND THE BOYBANDS

A range of proven techniques, from the careful and measured use of randomness to evolution, have been applied to generating music. Their use poses an interesting question: using software to create hit records seems possible, but is it desirable?

Rather than such systems creating fresh and exciting musical directions for us to explore, might they simply be employed by record companies to produce a never-ending stream of familiar-sounding singles, possibly guided directly by public feedback? Such software could, after all, toil day and night for different markets, never growing weary or picking up expensive habits along the way. **CS**